

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

5

APPLICATION PAPERS

10

OF

IGOR GARRIEVICH MUTTIK

15

FOR

20

25

DETECTING MALICIOUS COMPUTER PROGRAM ACTIVITY
USING EXTERNAL PROGRAM CALLS WITH DYNAMIC RULE SETS



BACKGROUND OF THE INVENTION

Field of the Invention

5 This invention relates to the field of data processing systems. More particularly, this invention relates to the detection of malicious computer program activity, such as, for example, as the result of computer viruses, worms, Trojans and the like, using detected characteristics of external program calls.

10 Description of the Prior Art

It is known to provide anti-virus systems which examine computer files to be accessed to determine characteristics of those computer files known to be indicative of malware, such as computer viruses, worms, Trojans and the like. Such techniques are necessarily reactive in that a new piece of malware must first be released and 15 identified before the characteristics of that new piece of malware can be determined such that the detection mechanisms can be updated to detect that new piece of malware. An alternative type of anti-virus system uses heuristic methods in which the computer file is examined to identify types of activity it is controlling and, if these match known suspicious types of activity, then an anti-malware response may be 20 triggered. Whilst such an approach can provide some protection against newly released malware, it does suffer from requiring a disadvantageously high processing overhead.

It is also known to provide anti-virus systems which examine the behaviour of 25 a computer system at the level of external program calls being made to the operating system. It has been found that certain sequences of external program calls, or combinations of external program calls with certain characteristics, are indicative of malicious computer program activity and may be used to trigger anti-malware responses. A problem with this behavioural analysis approach is that the malicious 30 activity against which a defence is being sought is allowed to occur prior to it being detected. Accordingly, some damage may already have been inflicted upon the computer system concerned by the malicious computer program activity, e.g. the malicious computer program may already have managed to replicate itself and infect further machines.

SUMMARY OF THE INVENTION

Viewed from one aspect the present invention provides a computer program product operable to detect malicious computer program activity, comprising:

- 5 logging code operable to log a stream of external program calls;
- primary set identifying code operable to identify, within said stream of external program calls, a primary set of one or more external program calls matching one or more rules indicative of malicious computer program activity from among a set of rules;
- 10 secondary set identifying code operable to identify, within said stream, at least one secondary set of one or more external program calls associated with said primary set of one or more external program calls; and
- modifying code operable to modify said set of rules such that said at least one secondary set of one or more external program calls are more strongly associated with malicious computer program activity.
- 15

The present technique recognises that external program calls logged in association with the primary set of external program calls known to correspond to malicious computer program activity may themselves subsequently be used as an indicator for malicious computer program activity. The secondary sets of external program calls are "tainted" by their association with the primary set of external program calls and the set of rules may be modified to be more sensitive to the secondary set of external program calls. In this way, the set of rules associated with malicious computer program activity may be extended and the detection made potentially more sensitive, reliable and proactive. It will be appreciated that in the context of the present technique, external program calls also encompasses calls to the standard libraries within the same program.

It will be appreciated that the secondary set of one or more external program calls could take any relative temporal location compared to the primary set of external program calls providing it is associated therewith, e.g. originates from the same computer program, memory region, thread etc, but in preferred embodiments the secondary set of one or more external program calls precedes the primary set of one or more external program calls within the stream of execution. In this way, the

modified/new rules added to this set of rules are able to detect malicious computer program activity earlier and possibly reduce or prevent damage caused by the malicious computer program activity.

5 It will be appreciated that the external computer program calls may take a variety of different forms, but these are preferably application program interface calls to an operating system. Such calls are clearly defined and accessible within a computer system. This makes them relatively easy to intercept and analyse. It will be appreciated that the computer calls being made can in themselves be analysed as well
10 as the results returned from those external program calls. The characteristics indicative of malicious computer activity may be characteristics of the returned results associated with the external program calls rather than the external program calls themselves or both.

15 The characteristics of the external program calls which may be analysed can vary considerably. In preferred embodiments these characteristics include at least one of a call name, a return address and one or more parameter values either outbound or inbound.

20 The rules against which the logged external program calls can be compared can take a wide variety of different forms, as will be familiar to those in the anti-malware field. A particularly convenient way of structuring this rule set is to use score values associated with certain characteristics or combinations of characteristics of external program calls. In this way, a logged stream may be analysed and when the
25 total score value associated with that logged stream exceeds a threshold value, an anti-malware response may be triggered. Pattern matching or regular expressions are other convenient ways of storing rules.

30 In the context of such a rule set using score values, a particularly convenient way of modifying the rule set make it more sensitive to the secondary sets of external program calls is to increase the score values associated with such secondary sets of external program calls.

It will be appreciated that the rule set or rule sets being used may advantageously be responsive in desired circumstances to the particular ordering of external program calls. It may be that with one ordering a set of external program calls has no malicious result whereas in another specific ordering a malicious result is
5 highly indicated.

The modification of the rule set can be seen effectively to perform a dynamic adaption of the set of rules in response to the detected stream of external program calls. Such a dynamic adaption provides a flexible approach to malware protection
10 and also provides the computer system with a degree of auto immunity whereby it may respond without programmer intervention in a manner which makes it more resistant to malware attack.

The modified set of rules may be kept locally and applied locally after its
15 modification and also may be advantageously transmitted to one or more remote computers and/or a rule supplier such that other users may also benefit from the modified rules.

It will be appreciated that the external program calls may be logged as they are
20 executed, but in a higher security system it may be desirable to emulate execution of a computer program to determine its external program calls which may then be logged and examined against the set of rules.

The modification of these set of rules can take a variety of different forms,
25 such as the above mentioned adjusting of the score values associated with certain sets and sequences of external program calls. In addition, new rules may be added to the existing set of rules as part of the modification of the set of rules.

In determining the secondary set of external program calls which may be
30 associated with the primary set of external program calls for which malicious computer program activity has been detected, it is advantageous to identify a starting point of the malicious computer program activity from which external program calls may be logged and examined. Such starting points may correspond to the start of

execution of a particular computer file, or a switch of memory region from which program instructions are being executed (e.g. a context switch).

In determining the validity and/or quality of the secondary set it is
5 advantageous to check if the modifications have positive effect. This can be decided
either internally (by applying some higher-level rules to the set modifications) or by
external signal(s). An example of such an external signal could be a report that
modified rule(s) decrease the malicious network traffic or slowdown the malware
propagation. As an example, after a modified set is transmitted to other computers
10 some network sensors detect the effect (e.g., decrease of traffic) and send a "positive"
signal back. That raises the score or promotes a rule from "temporary" into
"permanent" set.

Viewed from a further aspect a method of detecting malicious computer
15 program activity, said method comprising the steps of:

- logging a stream of external program calls;
- identifying within said stream of external program calls a primary set of one or
more external program calls matching one or more rules indicative of malicious
computer program activity from among a set of rules;
- 20 identifying within said stream at least one secondary set of one or more
external program calls associated with said primary set of one or more external
program calls; and
- modifying said set of rules such that said at least one secondary set of one or
more external program calls are more strongly associated with malicious computer
25 program activity.

Viewed from a still further aspect a data processing apparatus operable to
detect malicious computer program activity, said apparatus comprising:

- logging logic operable to log a stream of external program calls;
- 30 primary set identifying logic operable to identify, within said stream of
external program calls, a primary set of one or more external program calls matching
one or more rules indicative of malicious computer program activity from among a set
of rules;

secondary set identifying logic operable to identify, within said stream, at least one secondary set of one or more external program calls associated with said primary set of one or more external program calls; and

5 modifying logic operable to modify said set of rules such that said at least one secondary set of one or more external program calls are more strongly associated with malicious computer program activity.

10 The above, and other objects, features and advantages of this invention will be apparent from the following detailed description of illustrative embodiments which is to be read in connection with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 schematically illustrates a malware detecting system;

15 Figure 2 schematically illustrates the logging of a stream of external program calls and the examination of that stream to determine new detection rules;

20 Figure 3 is a flow diagram schematically illustrating processing steps associated with the examination of a stream of external program calls and the dynamic adaption of a set of rules for detecting malicious activity associated with that stream of external program calls;

Figure 4 schematically illustrates a modification to a set of rules; and

25 Figure 5 schematically illustrates the architecture of a general purpose computer which may be used to implement the above described techniques.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 illustrates a system for detecting malicious computer program activity.

30 An application program 2 executes in association with an operating system 4. The application program 2 issues external program calls, API calls, to the operating system 4 as is known and conventional in the computer programming field. These external calls result in desired processing operations being performed on behalf of the application program 2 and in some cases determined results are returned to the application program

2. The application program 2 may also communicate with a computer network 6 with which it may exchange data, messages, programs and the like.

As shown in Figure 1, an anti-malware layer 8 is provided between the application program 2 and the operating system 4. The anti-malware layer 8 serves to intercept external program calls issued by the application program 2 to the operating system 4 and subject these external program calls to checking by a set of rules 10. This checking is performed by external program call blocking code 12 which includes within its functionality the logging of a stream of external program calls, the identification of a primary set of program instruction calls found to match a rule or set of rules within the rules 10 and corresponding to malicious computer program activity, and thereafter identify associated secondary sets of external program calls within the log stream which should now be more strongly associated with malicious computer program activity. If the anti-malware layer 8 does not detect malicious computer program activity associated within the stream of external program calls, then these program calls are passed to the operating system 4 where they are serviced and the results, where applicable, returned back through the anti-malware layer 8 to the application program 2. The results associated with a particular external program call may also be examined and form part of the rule comparisons performed by the blocker code 12 in determining whether a particular external program call or set of external program calls matches one of the rules for identifying malicious computer program activity.

Figure 2 illustrates a stream of external program calls 14 which are logged by the anti-malware layer 8. A primary set of external program calls XYZ within this stream 14 is identified by one or more of the rules 10 as corresponding to malicious activity and will in its own right trigger an anti-malware response. As well as triggering such an anti-malware response, the anti-malware layer 8 also determines associated secondary sets of external program calls which are associated with the primary set XYZ. As will be seen, some of these secondary sets of external program calls precede the primary set XYZ and some succeed the primary set XYZ.

The secondary sets of external program calls which are selected to be used as potential new rules are selected from the preceding and succeeding external program calls as those external program calls known to be of most potential significance and

assistance in discriminating and identifying malicious computer activity. Some types of external program calls are common in both normal and malicious activity and accordingly of little use in selectively detecting malicious activity. However, other types of external program calls are unusual to an extent that when they occur having the 5 particular characteristics of those external program calls logged as preceding or succeeding the known malicious primary set of external program calls XYZ, then these are in themselves highly indicative of malicious computer program activity.

As illustrated in Figure 2, when the secondary sets of external program calls have 10 been so selected, new rules may be determined based upon them. These rules can require a certain ordering or spacing between specific external program calls or alternatively may detect the external program calls in any order providing they are within a certain distance of one another. The more highly specific the combination of external program calls identified in a new rule and known to be associated with the malicious 15 activity XYZ, the higher is made a score value to be associated with that new rule. A score value is used when a stream of external program calls is examined to determine whether or not a threshold value has been exceeded and anti-malware activity should be triggered.

20 In the example illustrated it will be seen that the primary set of external program calls XYZ which is detected as malicious activity results in the generation of plurality of new rules which serve to more strongly associate the secondary sets of external program calls with malicious activity. The secondary sets themselves may not be sufficient to trigger the anti-malware response, but their score values are increased such that when 25 they occur in combination with other detected behavioural characteristics an anti-malware response will now be triggered. The rules in themselves may not be new, it may rather be that the score values associated with a particular secondary set of external program calls is increased due to its now known association with the primary set XYZ of external program calls which is malicious.

30

Figure 3 is a flow diagram schematically illustrating the type of operation discussed in relation to Figure 2. At step 20 the anti-malware layer 8 waits for an API call to be received. At step 22 a received API call is added to a log. At step 24 the logged API calls (optionally in combination with their associated returned results) are

compared to a set of rules 10. If the result of the comparison with the rules 10 indicates malicious activity at step 26, then processing proceeds to step 28 where the API call concerned is blocked and an anti-malware action triggered. In case of a blocked API call(s) subsequent replies may be simulated to enable 'as-normal' further execution of the program. At step 30 the future stream of external program calls subsequent to the malicious activity is also logged (e.g. through emulation), such that these can be associated within the set of rules with the malicious activity as well as the already logged proceeding external program calls. At step 32, the logged external program calls are examined to determine secondary sets of external program calls which are to be subject to the generation of new rules or the modification of existing rules. In this review of the logged external program calls a starting point associated with the malicious computer activity is identified and used as a cut off proceeding which external program calls are not examined. Examples of such a starting point are the detected start of execution of a particular computer program and/or a context switch detected by a change in the memory region from which a computer program instructions are being fetched.

Once the secondary sets of external program calls to have new or modified rules associated with them have been identified at step 32, the new rules concerned are generated at step 34 and updated into the set of rules 10 at step 36. The new rules or the complete new set of rules are then transmitted to a central rules supplier and optionally other remote computers at step 38 (via a network link 6), such that the supplier and the remote computers can benefit from the dynamic adaptation of the rules 10 which has occurred. Processing then returns to step 20.

If the determination at step 26 was that malicious activity was not associated with the newly-received external program call, then processing proceeds to step 40 at which that external program call is passed to the operating system 4 in the normal way.

Figure 4 schematically illustrates a set of rules 10 and a modified set of rules 10'. As illustrated, both the original set of rules 10 and the modified set of rules 10' include a rule identifying the primary set of external program calls XYZ with a score value of 110. This is above the threshold value of 100 taken in this example to be indicative of malicious computer program activity. In the modified set of rules 10', the sequences of secondary sets of external program calls illustrated in Figure 2 have been added together

with their associated score values. The secondary set of external program calls AB was already present within the original set of rules 10, but in this case has had its score value increased from 50 to 70. New rules associated with the other four examples given in Figure 2 have been appended to the end of the modified set of rules 10' and will in future 5 be applied in detecting malicious computer activity by examining the logged stream of external program calls.

Figure 5 schematically illustrates a general purpose computer 200 of the type that may be used to implement the above described techniques. The general purpose 10 computer 200 includes a central processing unit 202, a random access memory 204, a read only memory 206, a network interface card 208, a hard disk drive 210, a display driver 212 and monitor 214 and a user input/output circuit 216 with a keyboard 218 and mouse 220 all connected via a common bus 222. In operation the central processing unit 202 will execute computer program instructions that may be stored in one or more of the 15 random access memory 204, the read only memory 206 and the hard disk drive 210 or dynamically downloaded via the network interface card 208. The results of the processing performed may be displayed to a user via the display driver 212 and the monitor 214. User inputs for controlling the operation of the general purpose computer 200 may be received via the user input output circuit 216 from the keyboard 218 or the 20 mouse 220. It will be appreciated that the computer program could be written in a variety of different computer languages. The computer program may be stored and distributed on a recording medium or dynamically downloaded to the general purpose computer 200. When operating under control of an appropriate computer program, the 25 general purpose computer 200 can perform the above described techniques and can be considered to form an apparatus for performing the above described technique. The architecture of the general purpose computer 200 could vary considerably and Figure 6 is only one example.

Although illustrative embodiments of the invention have been described in detail 30 herein with reference to the accompanying drawings, it is to be understood that the invention is not limited to those precise embodiments, and that various changes and modifications can be effected therein by one skilled in the art without departing from the scope and spirit of the invention as defined by the appended claims.